



EmoSphere-TTS: Emotional Style and Intensity Modeling via Spherical Emotion Vector for Controllable Emotional Text-to-Speech

Deok-Hyeon Cho¹, Hyung-Seok Oh¹, Seung-Bin Kim¹, Sang-Hoon Lee², Seong-Whan Lee^{†1}

¹Department of Artificial Intelligence, Korea University, Seoul, Korea

²Department of Software and Computer Engineering, Ajou University, Gyeonggi-do, Korea

dh_cho@korea.ac.kr, hs_oh@korea.ac.kr, sb-kim@korea.ac.kr, sanghoonlee@ajou.ac.kr, sw.lee@korea.ac.kr

Abstract

Despite rapid advances in the field of emotional text-to-speech (TTS), recent studies primarily focus on mimicking the average style of a particular emotion. As a result, the ability to manipulate speech emotion remains constrained to several predefined labels, compromising the ability to reflect the nuanced variations of emotion. In this paper, we propose EmoSphere-TTS, which synthesizes expressive emotional speech by using a spherical emotion vector to control the emotional style and intensity of the synthetic speech. Without any human annotation, we use the arousal, valence, and dominance pseudo-labels to model the complex nature of emotion via a Cartesian-spherical transformation. Furthermore, we propose a dual conditional adversarial network to improve the quality of generated speech by reflecting the multi-aspect characteristics. The experimental results demonstrate the model's ability to control emotional style and intensity with high-quality expressive speech.

Index Terms: Text-to-speech, expressive emotional speech synthesis, emotional style and intensity control

1. Introduction

Recently, text-to-speech (TTS) has shown rapid progress [1, 2, 3], and interest in emotional TTS has also increased. Although emotional TTS technology has seen significant improvements in recent years, there remains a limitation in high-level interpretable emotion control [4, 5, 6, 7]. In particular, emotional control remains challenging, as speech labeled with the same emotion can exhibit diverse emotional expressions greatly influenced by the variability in acting performances.

For emotional TTS, a common approach is to control the diverse emotional expressions from the emotion labels and the reference audio. The emotion label-based approach models the complex nature of emotion to serve as an auxiliary input for the TTS system. Relative attribute [8, 9, 10] is one of the most representative methods, utilizing a learned ranking function [11] to delineate differences between binary classes. EmoQ-TTS [12] employs distance-based quantization via linear discriminant analysis to control fine-grained emotional intensity. Another way to control the diverse expression of emotion is through reference-based emotional TTS. In this context, a scaling factor [13, 14] is multiplied by the emotion embedding to control emotion intensity precisely. However, these methods have several limitations. All methods based on emotion labels focus on transferring emotions using discrete labels that ignore the complex nature of emotion conveyed in human speech [15]. For example, while *sad* spans emotions like *lonely* and *hurt*, its categorization typically reduces expressions to a uni-

form style. Using references to control emotional expression is further complicated by the difficulty in capturing the nuances of references due to a mismatch between the reference and the synthesized speech. Furthermore, these methods face difficulty finding suitable scaling factors, and making adjustments often results in unstable audio quality.

Another strategy to control emotional expression involves leveraging the emotional dimensions. Russell [16] proposed a continuous emotion space as an alternative for human emotions. Building on this, researchers attempt to control emotional expression by utilizing the extended emotional dimensions of arousal, valence, and dominance (AVD) [17, 18]. Emotional dimensions provide a continuous and fine-grained description, offering more detailed control than discrete emotions. However, only a few emotional speech databases provide these annotations due to the inherent subjectivity and the high costs associated with collecting such data. Furthermore, emotional dimensions are challenging for models to control intuitively and distinctly characterize the diverse emotional expression.

To address the above limitations, we propose EmoSphere-TTS, a novel approach to control emotional style and intensity with spherical emotion vector space in emotional TTS. We adopt the emotional dimensions of AVD from pseudo-labeling in speech emotion recognition (SER). We also propose a spherical emotion vector space via Cartesian-spherical transformation to model the complex nature of emotion that has been difficult in Cartesian coordinate systems. We found that this space is the key to controlling the emotional style and intensity of the synthetic speech. Furthermore, we introduce dual conditional adversarial training to improve the quality of generated speech by reflecting the emotion and speaker-specific characteristics. The experimental results demonstrate that our model outperforms the others in terms of audio quality and emotion similarity on the controllable emotional TTS. Audio samples are available at <https://EmoSphere-TTS.github.io/>.

2. EmoSphere-TTS

We present a controllable emotional TTS system, EmoSphere-TTS. We introduce spherical emotion vector space and spherical emotion encoder to deliver speech with the complex nature of emotion. Furthermore, we introduce a dual conditional discriminator for better audio quality and expressiveness. The details are described in the following subsections.

2.1. Emotional style and intensity modeling

In this section, we model the diverse emotional expressions through a spherical emotion vector space. The approach to building the space is structured around two key components: i) the AVD encoder and ii) the Cartesian-spherical transformation.

[†]Corresponding author

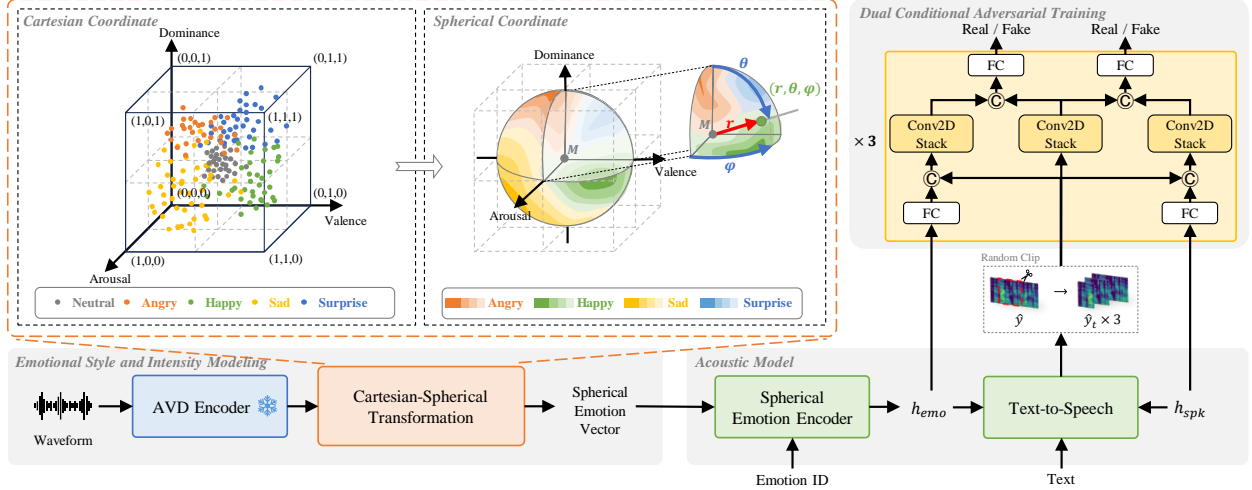


Figure 1: Overall architecture of EmoSphere-TTS.

2.1.1. AVD encoder

Instead of using emotional dimensions of human annotation, we adopt wav2vec 2.0 [19]-based SER [20] to extract consistently continuous and detailed representations from audio. The model generates predictions for $e_{ki} = (d_a, d_v, d_d)$, where d_a represents arousal, d_v valence, and d_d dominance, each ranging approximately from 0 to 1 in Cartesian coordinates. Here, e_{ki} denotes the i -th coordinate of the k -th emotion.

2.1.2. Cartesian-spherical transformation

For modeling the complex nature of emotion, we introduce the spherical emotion vector space, which represents the relative distance and angle vector from the neutral center. Inspired by coordinate transformations [21], which can be easily controlled by a continuous scalar indicating emotional style and intensity. We transform all points for AVD pseudo-labels in spherical coordinates by following the assumptions: i) the emotional intensity increases as it moves farther from the neutral emotion center, and ii) the angle from the neutral emotion center determines the emotional style. First, we obtain transformed Cartesian coordinates $e'_{ki} = (d'_a, d'_v, d'_d)$ by setting the neutral emotion center M as the origin,

$$e'_{ki} = e_{ki} - M \text{ where } M = \frac{1}{N_n} \sum_{i=1}^{N_n} e_{ni}, \quad (1)$$

where N_n represents the total number of neutral coordinates e_{ni} . Then, the transformation from Cartesian coordinates to spherical coordinate (r, ϑ, φ) can be formulated as:

$$r = \sqrt{d'_a{}^2 + d'_v{}^2 + d'_d{}^2}, \quad (2)$$

$$\vartheta = \arccos\left(\frac{d'_d}{r}\right), \varphi = \arctan\left(\frac{d'_v}{d'_a}\right). \quad (3)$$

After the Cartesian-spherical transformation, we normalize the intensity of the emotion by scaling the radial distance r to a range of 0 to 1. To achieve this, the min-max normalization process utilizes the interquartile range technique [22], effectively determining the minimum and maximum values for the scale. Additionally, we quantize the emotion style by segmenting directional angles ϑ and φ into eight octants, each defined by the positive and negative directions of the A , V , and D axes.

2.2. Spherical emotion encoder

After building spherical emotion vector space, the spherical emotion encoder blends them with emotion ID to compose their spherical emotion embedding. Initially, we use a projection layer to align the dimensions of the emotion style vector and emotion class embedding. Then, we concatenate these projections and apply a softplus activation [25] similar to [26] followed by Layer Normalization (LN) [27]. Finally, the spherical emotion embedding, \mathbf{h}_{emo} , is merged with projected emotion intensity vectors, as the following equation:

$$\mathbf{h}_{emo} = \text{LN}(\text{softplus}(\text{concat}(\mathbf{h}_{sty}, \mathbf{h}_{cls}))) + \mathbf{h}_{int}. \quad (4)$$

Here, \mathbf{h}_{sty} , \mathbf{h}_{int} , and \mathbf{h}_{cls} denote the outputs of the projection layer related to the emotional style vector, emotional intensity vector, and emotion class embedding, respectively.

2.3. Dual conditional adversarial training

We adopt the structure of multiple CNN-based discriminators [28] for adversarial training to improve the quality of the TTS model. These discriminators comprise a Conv2D stack that consists of multiple stacked 2D-convolutional layers and fully connected (FC) layers. The input value is the random Mel-spectrogram clip (Mel clip) using random windows of different lengths t . To improve quality and further expressiveness, we utilize emotion and speaker embeddings to capture the multi-aspect characteristics more effectively, inspired by [29, 30]. One Conv2D stack receives only the Mel clip, while the others receive a combination of condition embedding and the Mel clip. We extend the condition embedding to match the length of the Mel clip for concatenation. The loss functions \mathcal{L} of the discriminator D and generator G are shown in Equation (5) and (6):

$$\mathcal{L}_D = \sum_{c \in \{spk, emo\}} \sum_t \mathbb{E}[(1 - D_t(y_t, c))^2 + D_t(\hat{y}_t, c)^2], \quad (5)$$

$$\mathcal{L}_G = \sum_{c \in \{spk, emo\}} \sum_t \mathbb{E}[(1 - D_t(\hat{y}_t, c))^2], \quad (6)$$

where y_t and \hat{y}_t respectively represent the ground truth and generated Mel-spectrograms, with c denoting the condition type.

Table 1: Subjective and objective evaluation results. The last group of rows represents an ablation study of the proposed model. The nMOS and sMOS scores are presented with 95% confidence intervals.

Method	UTMOS (\uparrow)	nMOS (\uparrow)	sMOS (\uparrow)	WER (\downarrow)	CER (\downarrow)	RMSE _{f₀} (\downarrow)	RMSE _{period} (\downarrow)	F1 V/UV (\uparrow)	EER (\downarrow)	SECS (\uparrow)	ECA (\uparrow)
GT	3.78	4.27±0.04	4.37±0.11	11.92	3.04	-	-	-	2.99	0.753	85.67
BigVGAN [23]	3.62	4.26±0.04	4.30±0.11	11.97	3.03	32.21	0.214	0.8630	3.19	0.741	86.07
FastSpeech 2 w/ Emotion Label [24]	2.52	3.13±0.06	3.28±0.15	21.09	9.30	44.65	0.455	0.6995	5.89	0.654	93.75
FastSpeech 2 w/ Relative Attribute [9]	2.57	3.16±0.06	3.22±0.14	23.70	10.02	44.82	0.454	0.6978	6.20	0.654	92.80
FastSpeech 2 w/ Scaling Factor [14]	2.66	3.34±0.05	3.29±0.13	19.46	8.73	41.90	0.453	0.7028	6.39	0.645	64.19
EmoSphere-TTS (Proposed)	3.15	3.88±0.05	3.48±0.11	17.43	7.05	42.57	0.447	0.7077	4.29	0.669	94.02
w/o Spherical Emotion Vector	3.04	3.69±0.05	3.45±0.12	19.22	8.05	43.70	0.449	0.7099	5.50	0.649	93.89
w/o Dual Conditional Discriminator	2.86	3.39±0.06	3.24±0.14	18.09	7.38	42.52	0.449	0.7055	5.86	0.666	92.60

2.4. TTS model

We retain the original architecture and objective function of FastSpeech 2 [24] except for using an emotion spherical vector to provide emotional style and intensity information. Additionally, the speaker ID is mapped into an embedding h_{spk} to represent different speaker characteristics. Then, the speaker and emotion embedding are concatenated and provided to the variance adaptor. During inference, we use manual style and intensity vectors to control the diverse emotional expressions. By manipulating the emotional style and intensity in the spherical emotion vector space, we can efficiently synthesize the complex nature of emotion and control the diverse emotional expressions in synthesized speech.

3. Experiments and results

3.1. Experimental setup

We use the emotional speech dataset (ESD) [31], which consists of about 350 parallel utterances spoken by 10 English speakers with five emotional states (neutral, happy, angry, sad, and surprise). Following the prescribed data partitioning criteria, we extracted one sample for each emotion from every speaker, resulting in a total of 17,500 samples. For the Mel-spectrogram, we transform audio using the short-time Fourier transform with a hop size of 256, a window size of 1,024, an FFT size of 1,024, and 80 bins of Mel-filter. We employ the AdamW optimizer [32], setting the hyperparameters β_1 to 0.9 and β_2 to 0.98. For the training of the TTS system and discriminator, the learning rates were configured at 5×10^{-4} and 1×10^{-4} , respectively. The training process of the TTS module was conducted over approximately 24 hours on a single NVIDIA RTX A6000 GPU. For the audio synthesis in our experiments, we utilize the official implementation of BigVGAN [23], along with its pre-trained model.

3.2. Implementation details

For the acoustic model, following the FastSpeech 2 [24] configuration, in the FFT block of the phoneme encoder and decoder, we configure the number of layers to 4, hidden size to 256, filter size to 1024, and kernel size to 9. Regarding the AVD encoder, we adopt a system proposed in [20], which predicts AVD using wav2vec 2.0 [19] and a linear predictor. For the discriminator, we use projection layers for each condition with a hidden size of 128 and three different sizes of windows ([32, 64, 96]).

3.3. Model performance

We compare EmoSphere-TTS with other systems: 1) FastSpeech 2 w/ emotion label [24], which adopts emotion id with a look-up table. 2) FastSpeech 2 w/ relative attribute [9], which assigns emotion intensity via learned ranking function [11]. 3) FastSpeech 2 w/ scaling factor [14], which is multiplied by the emotion embedding to control emotion intensity precisely. We

Table 2: Results of evaluation for emotion intensity control. Emotion Intensity Recognition is asked to select the sample with the stronger intensity from a pair. Weak to Strong represents emotional intensity.

Emotion	Method	Emotion Intensity Recognition [%]		
		Weak<Medium	Medium<Strong	Weak<Strong
Angry	Relative Attribute [9]	0.56	0.56	0.72
	Scaling Factor [14]	0.44	0.52	0.48
	EmoSphere-TTS	0.72	0.75	0.79
Sad	Relative Attribute [9]	0.57	0.39	0.55
	Scaling Factor [14]	0.69	0.48	0.63
	EmoSphere-TTS	0.62	0.42	0.48
Happy	Relative Attribute [9]	0.48	0.51	0.60
	Scaling Factor [14]	0.61	0.43	0.44
	EmoSphere-TTS	0.80	0.66	0.84
Surprise	Relative Attribute [9]	0.44	0.55	0.52
	Scaling Factor [14]	0.50	0.45	0.46
	EmoSphere-TTS	0.69	0.76	0.79
Average	Relative Attribute [9]	0.52	0.51	0.60
	Scaling Factor [14]	0.56	0.47	0.50
	EmoSphere-TTS	0.71	0.65	0.72

follow the same experimental setup and model configurations to ensure a fair comparison. For evaluation metrics, we use a naturalness mean opinion score (nMOS) and a similarity mean opinion score (sMOS). The nMOS and sMOS are reported with 95% confidence intervals. Additionally, we utilize the open-source UTMOS¹ [33] as a MOS prediction model for the naturalness metric. To evaluate linguistic consistency, we calculate the word error rate (WER) and character error rate (CER) by Whisper large model [34]. For the speaker similarity measurements, we calculate the speaker embedding cosine similarity (SECS) via Resemblyzer² between the target and converted speech and equal error rate (EER) via the automatic speech recognition model [35]. For emotionally expressive evaluation, we conduct emotion classification accuracy (ECA) using a pre-built emotion classification model [36]. For prosodic evaluation, we compute the root mean square error for both pitch error (RMSE_{f₀}) and periodicity error (RMSE_{period}), along with the F1 score of voiced/unvoiced classification (F1_{v/uv}).

Given our primary on quality and expressiveness, manual intensity and style vectors are not employed in these experiments. In the ablation study, w/o spherical emotion vector is a model that uses emotion id with a lookup table instead of a spherical emotion vector and encoder. As shown in Table 1, our model achieves significant improvements, and this can be explained by: 1) as opposed to transferring emotion from emotion label-based or reference audio methods, directly assigning spherical emotion vector is easier for the model to generate good quality speech. Our spherical emotion vector exhibits better expressiveness and audio quality, including naturalness and pronunciation, even without a dual conditional discriminator. 2) The dual conditional discriminator improves the quality of the generated speech by reflecting both emotion and speaker characteristics.

¹<https://github.com/tarepan/SpeechMOS>

²<https://github.com/resemble-ai/Resemblyzer>




Figure 2: The pitch tendency according to emotion and intensity.

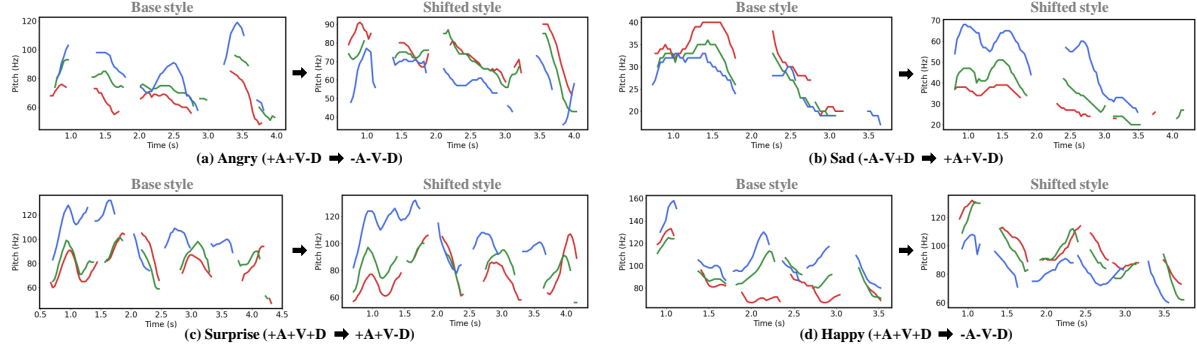


Figure 3: A pitch track of a sample demonstrating the effects of emotional style shift, where A , V , and D represent arousal, valence, and dominance, respectively. The line color represents emotional intensity, red = 0.1, green = 0.5, and blue = 0.9.

3.4. Emotion intensity controllability

In this section, we conduct a subjective evaluation to determine the discernibility of synthesized speech samples exhibiting varying intensity levels. To demonstrate the intensity control capability of our model, we synthesize speech with three different levels of emotion intensity (weak, medium, and strong). Evaluators are presented with two different sentences, each with varying intensities, and tasked with selecting the one that exhibits stronger emotion. In relative attribute [9] and EmoSphere-TTS, we uniformly refer to scores 0.1 as weak, 0.5 as medium, and 0.9 as strong. Scaling factor [14] cannot assign intensity scores, so we set the scalar factor to 1, 2, and 3 to represent the weak, medium, and strong emotion strength the same as in the original setting. Additionally, we visualize the tendency of pitch to demonstrate the ability to control emotional expression as shown in Figure 2. We computed these values by averaging the pitch from the synthesized speeches by all combinations of emotion labels and intensity vectors for all test sentences.

As shown in Table 2, relative attribute [9] effectively controls the intensity. However, in the *sad* emotional speech, the pitch increases as the intensity increases, as shown in Figure 2 (a). This indicates that subtle emotional nuances are complex to capture when considering emotion labels alone, and expressions are often reduced to a uniform style. Scaling factor [14] might not be efficient for performing intensity control; in some cases, the intensity difference is not readily perceivable, as shown in Figure 2 (b). However, as shown in Table 2, the scaling factor outperformed the other models for the *sad* emotion. Still, the scaling factor focuses on reducing pitch and decelerating speech rate in static emotions, overlooking the complex nature of emotion. On the other hand, compared to the baseline models, EmoSphere-TTS performs the best. Furthermore, the pitch tendency plot reflects the intensity according to the emotion. This indicates that the proposed model synthesizes speech according to a given intensity scale.

3.5. Emotion style shift

To demonstrate the changing patterns of emotion intensity based on the shifted emotion style, we visualize the pitch tracks of a sample. In Figure 3, we observe that when the base style vector is input, emotion intensity changing patterns reflect the characteristics of the AVD axes. For example, style vectors with positive A axes have a pitch that tends to increase the changing patterns, positive V axes have higher average pitch values, and positive D axes have a narrow range in changing patterns. This indicates that arousal, valence, and dominance reflect each meaning of axes, representing the level of excitement or energy, positivity or negativity of emotion, and control level within an emotional state, respectively. By shifting the style vector, the emotion intensity patterns change with the shifted axis. This indicates that the proposed spherical emotion vector reflects diverse emotional expressions and offers detailed manipulation of emotional expressions.

4. Conclusion

We present EmoSphere-TTS, a system that synthesizes expressive emotional speech through a spherical emotion vector space, controlling the diverse emotion expression. With only a speech dataset, we extract AVD pseudo-labels and model generalized representations of the emotional style and intensity through Cartesian-spherical transformation. Furthermore, we improve the quality and emotional expressiveness of the overall model using the dual conditional adversarial discriminator and spherical emotion encoder. The experimental results demonstrate that our proposed spherical emotion vector effectively synthesizes the complex nature of emotion and controls diverse emotion expression. In this article, we only focused exclusively on the global style present within sentence-level emotional information. In future work, we aim to extend our approach to include phoneme-level emotional information and allow for fine-grained control. We also expect that the proposed method can be utilized for emotional voice conversion like [36, 37, 38].

5. Acknowledgements

This work was partly supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program (Korea University), No. 2021-0-02068, Artificial Intelligence Innovation Hub, and AI Technology for Interactive Communication of Language Impaired Individuals).

6. References

- [1] Y. Wang, R. Skerry-Ryan, D. Stanton, Y. Wu, R. J. Weiss, N. Jaitly, Z. Yang, Y. Xiao, Z. Chen, S. Bengio, Q. Le, Y. Agiomyriannakis, R. Clark, and R. A. Saurous, "Tacotron: Towards End-to-End Speech Synthesis," in *Proc. INTERSPEECH 2017*, 2017, pp. 4006–4010.
- [2] S.-H. Lee, S.-B. Kim, J.-H. Lee, E. Song, M.-J. Hwang, and S.-W. Lee, "Hierspeech: Bridging the gap between text and speech by hierarchical variational inference using self-supervised representations for speech synthesis," *Proc. Adv. Neural Inf. Process. Syst.*, 2022.
- [3] S.-B. Kim, S.-H. Lee, H.-Y. Choi, and S.-W. Lee, "Audio super-resolution with robust speech representation learning of masked autoencoder," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 32, pp. 1012–1022, 2024.
- [4] S. Wang, J. Guonason, and D. Borth, "Learning Emotional Representations from Imbalanced Speech Data for Speech Emotion Recognition and Emotional Text-to-Speech," in *Proc. INTERSPEECH 2023*, 2023.
- [5] G. Liu, Y. Zhang, Y. Lei, Y. Chen, R. Wang, L. Xie, and Z. Li, "PromptStyle: Controllable Style Transfer for Text-to-Speech with Natural Language Descriptions," in *Proc. INTERSPEECH 2023*, 2023.
- [6] R. Liu, H. Zuo, D. Hu, G. Gao, and H. Li, "Explicit Intensity Control for Accented Text-to-speech," in *Proc. INTERSPEECH 2023*, 2023.
- [7] M. Kang, W. Han, S. J. Hwang, and E. Yang, "ZET-Speech: Zero-shot adaptive Emotion-controllable Text-to-Speech Synthesis with Diffusion and Style-based Models," in *Proc. INTERSPEECH 2023*, 2023.
- [8] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, "Speech synthesis with mixed emotions," *Proc. IEEE Trans. Affect. Comput.*, vol. 14, no. 4, pp. 3120–3134, 2023.
- [9] X. Zhu, S. Yang, G. Yang, and L. Xie, "Controlling emotion strength with relative attribute for end-to-end speech synthesis," in *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2019, pp. 192–199.
- [10] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, "Emotion intensity and its control for emotional voice conversion," *Proc. IEEE Trans. Affect. Comput.*, vol. 14, no. 1, pp. 31–48, 2022.
- [11] D. Parikh and K. Grauman, "Relative attributes," in *Proc. IEEE Int. Conf. Comput. Vis.* IEEE, 2011.
- [12] C.-B. Im, S.-H. Lee, S.-B. Kim, and S.-W. Lee, "Emoq-tts: Emotion intensity quantization for fine-grained controllable emotional text-to-speech," in *IEEE Int. Conf. Acoust., Speech Signal Process.* IEEE, 2022.
- [13] T. Li, X. Wang, Q. Xie, Z. Wang, M. Jiang, and L. Xie, "Cross-speaker Emotion Transfer Based On Prosody Compensation for End-to-End Speech Synthesis," in *Proc. INTERSPEECH 2022*, 2022, pp. 5498–5502.
- [14] T. Li, X. Wang, Q. Xie, Z. Wang, and L. Xie, "Cross-speaker emotion disentangling and transfer for end-to-end speech synthesis," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 30, pp. 1448–1460, 2022.
- [15] Y. Shin, Y. Lee, S. Jo, Y. Hwang, and T. Kim, "Text-driven Emotional Style Control and Cross-speaker Style Transfer in Neural TTS," in *Proc. INTERSPEECH 2022*, 2022.
- [16] J. A. Russell, "A circumplex model of affect," *Journal of personality and social psychology*, vol. 39, no. 6, p. 1161, 1980.
- [17] R. Habib, S. Mariooryad, M. Shannon, E. Battenberg, R. Skerry-Ryan, D. Stanton, D. Kao, and T. Bagby, "Semi-supervised generative modeling for controllable speech synthesis," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [18] S. Sivaprasad, S. Kosgi, and V. Gandhi, "Emotional Prosody Control for Speech Generation," in *Proc. INTERSPEECH 2021*, 2021.
- [19] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Proc. Adv. Neural Inf. Process. Syst.*, 2020.
- [20] J. Wagner, A. Triantafyllopoulos, H. Wierstorf, M. Schmitt, F. Burkhardt, F. Eyben, and B. W. Schuller, "Dawn of the transformer era in speech emotion recognition: closing the valence gap," *IEEE Trans. Pattern Anal. Mach. Intell.*, 2023.
- [21] R. Jenke and A. Peer, "A cognitive architecture for modeling emotion dynamics: Intensity estimation from physiological signals," *Cognitive Systems Research*, vol. 49, pp. 128–141, 2018.
- [22] S. Walfish, "A review of statistical outlier methods," *Pharmaceutical technology*, vol. 30, no. 11, p. 82, 2006.
- [23] S.-g. Lee, W. Ping, B. Ginsburg, B. Catanzaro, and S. Yoon, "Bigvgan: A universal neural vocoder with large-scale training," in *Proc. Int. Conf. Learn. Representations*, 2023.
- [24] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [25] H. Zheng, Z. Yang, W. Liu, J. Liang, and Y. Li, "Improving deep neural networks using softplus units," in *Proc. Int. Joint Conf. Neural Netw.* IEEE, 2015, pp. 1–4.
- [26] H.-W. Yoon, O. Kwon, H. Lee, R. Yamamoto, E. Song, J.-M. Kim, and M.-J. Hwang, "Language Model-Based Emotion Prediction Methods for Emotional Speech Synthesis Systems," in *Proc. INTERSPEECH 2022*, 2022.
- [27] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [28] Z. Ye, Z. Zhao, Y. Ren, and F. Wu, "Syntaspeech: Syntax-aware generative adversarial text-to-speech," *Proc. Int. Joint Conf. Artif. Intell.*, 2022.
- [29] J. Yang, J.-S. Bae, T. Bak, Y.-I. Kim, and H.-Y. Cho, "GANSpeech: Adversarial Training for High-Fidelity Multi-Speaker Speech Synthesis," in *Proc. INTERSPEECH 2021*, 2021.
- [30] H.-S. Oh, S.-H. Lee, and S.-W. Lee, "Diffprosody: Diffusion-based latent prosody generation for expressive speech synthesis with prosody conditional adversarial training," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 32, pp. 2654–2666, 2024.
- [31] K. Zhou, B. Sisman, R. Liu, and H. Li, "Emotional voice conversion: Theory, databases and esd," *Speech Communication*, vol. 137, pp. 1–18, 2022.
- [32] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [33] T. Saeki, D. Xin, W. Nakata, T. Koriyama, S. Takamichi, and H. Saruwatari, "UTMOS: UTokyo-SaruLab System for Voice-MOS Challenge 2022," in *Proc. INTERSPEECH 2022*, 2022.
- [34] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *Int. Conf. Mach. Learn.*, 2023, pp. 28 492–28 518.
- [35] Y. Kwon, H.-S. Heo, B.-J. Lee, and J. S. Chung, "The ins and outs of speaker recognition: lessons from voxsrc 2020," in *IEEE Int. Conf. Acoust., Speech Signal Process.* IEEE, 2021.
- [36] K. Zhou, B. Sisman, R. Rana, B. W. Schuller, and H. Li, "Emotion intensity and its control for emotional voice conversion," *Proc. IEEE Trans. Affect. Comput.*, vol. 14, no. 1, pp. 31–48, 2023.
- [37] S.-H. Lee, H.-Y. Choi, H.-S. Oh, and S.-W. Lee, "HierVST: Hierarchical Adaptive Zero-shot Voice Style Transfer," in *Proc. INTERSPEECH 2023*, 2023.
- [38] H.-S. Oh, S.-H. Lee, D.-H. Cho, and S.-W. Lee, "Durflex-evc: Duration-flexible emotional voice conversion with parallel generation," *arXiv preprint arXiv:2401.08095*, 2024.